HITACHI

Best Practice for Hitachi Virtual Storage Platform with Hyper-V on Windows Failover Cluster

MK-SL-438-01
September 2026

Hitachi Vantara

© Hitachi Vantara LLC 2026. All Rights Reserved.

Table of Contents

Notices and Disclaimer 3

About This Document 4

Introduction 4
Intended Audience 4
Value Proposition 4
Document Revisions 5
Comments 5

Fibre Channel (FC) Configuration for Windows Environments 6

Host Side Requirements 6
FC Switch Zoning 6

Microsoft Windows Host Configuration with HV Storage 7

Target (Hitachi storage) Queue Depth 7
Initiator Queue Depth Settings 7
Host Group and LUN path configuration 8
Configuring Host mode and Host mode options 9

Install MPIO on Windows Server 2025 10

Accessing the MPIO Control Panel 10
Understanding the MPIO Properties Dialog Box 10
Configuring the parameter (Microsoft DSM/MPIO) 11

Windows Failover Cluster 12

Quorum Configuration 12
Node Failover and Failback Policy 12
Network Configuration 12
Validating the Cluster Configuration 13
Cluster Shared Volumes (CSV) 14

Windows Hyper-V on Failover Cluster 15

© Hitachi Vantara LLC 2026. All Rights Reserved.

Hyper-V Best Practices on VSP Block Storage 15
Recommended Configuration Guidelines 15
Hyper-V Networking Fundamentals 15
Basic Network Configuration 15
Processor and Memory Allocation 16
Best Practices for Choosing Hyper-V Virtual Machine Generation 17
Virtual Machines on Hitachi SAN Disks using Virtual SAN 17

Terms and definitions 21

© Hitachi Vantara LLC 2026. All Rights Reserved.

Notices and Disclaimer

© 2026 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval systemfor commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee maynot make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at anytime without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara at https://support.HitachiVantara.com/ en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara.

By using this software, you agree that you are responsible for:

1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, andLicensee shall not export, or allowthe export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., In the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPii™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

About This Document

Introduction

As organizations continue to modernize their datacenter infrastructure, the combination of Microsoft Windows Server Failover Clustering (WSFC) and Hyper-V remains a widely adopted solution for building highly available and scalable virtualization platforms. When paired with enterprise-grade storage solutions such as Hitachi Virtual Storage Platform (VSP), this architecture delivers a robust foundation capable of supporting mission-critical workloads with exceptional performance, reliability, and operational efficiency.

Deploying Hyper-V on a Windows Failover Cluster backed by VSP storage, however, requires careful planning and adherence to proven best practices. Proper configuration of cluster components, storage provisioning, networking, and failover behavior is essential to ensure seamless virtual machine availability, optimal I/O performance, and long-term maintainability. This document provides a comprehensive overview of the recommended best practices for designing and implementing this deployment model.

The following table details applicable support information.

Operating System	Hitachi Storage Models
Windows Server 2025 Windows Server 2022	VSP One Block 85 VSP One Block 20 Series VSP E Series VSP 5000 Series

Intended Audience

This document is intended for architects, storage administrators, system administrators, and system engineers with a thorough knowledge of using Hitachi Virtual Storage Platform One Block storage, Microsoft Windows Operating System, Virtualization, Fibre Channel working methodology, Fibre Channel switches, servers, and Fibre Channel HBA/Initiators to create a resilient, high-performance virtualized environment that leverages the full capabilities of Microsoft virtualization technologies and VSP storage systems.

Value Proposition

The Best Practice Guide for Optimized Performance for HyperV on Windows Server Failover Cluster using VSP storage empowers IT architects, virtualization engineers, and datacenter administrators to build a high-availability, high-performance virtualization platform with confidence. By combining Microsoft's proven clustering and HyperV capabilities with the enterprise class reliability of VSP storage, this guide delivers practical, actionable insights that enable organizations to:

- **Maximize Virtualization Performance**
  Ensure consistently high I/O throughput, low latency, and optimized VM operations through validated storage and cluster configuration practices.
- **Enhance Infrastructure Resilience and Availability**
  Implement a robust architecture that leverages Windows Server Failover Clustering and the advanced data services of VSP storage systems for seamless failover, fault tolerance, and workload continuity.
- **Simplify Deployment and Reduce Operational Risk**
  Follow prescriptive, field tested recommendations that streamline implementation, reduce misconfiguration, and minimize troubleshooting with time-tested recommendations that streamline implementation, reduce misconfiguration, and minimize troubleshooting time.
- **Achieve Predictable, Scalable Growth**
  Utilize a design framework that supports workload expansion, storage tiering, and performance scaling without compromising stability or efficiency.

© Hitachi Vantara LLC 2026. All Rights Reserved

Document Revisions

Revision Number	Date	Details
MK-SL-438-00	February 2026	Initial release
MK-SL-438-01	September 2026	Parameter updates

Comments

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to Docs-Feedback@hitachivantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Thank you!

© Hitachi Vantara LLC 2026. All Rights Reserved

Fibre Channel (FC) Configuration for Windows Environments

Hitachi Virtual Storage Platform (VSP) storage systems fully support traditional Fibre Channel (FC-SCSI) connectivity with Windows Server. For details on supported Windows versions, FC HBAs, and VSP storage models, refer to the Hitachi Interoperability Matrix.

Windows Server provides native support for FC environments using standard Fibre Channel Host Bus Adapters (HBAs). FC is used to present block storage (LUNs) directly from the VSP storage system to the Windows host over an FC SAN.

Host Side Requirements

FC HBA Support

Ensure the Windows host is equipped with the following:

- A supported Fibre Channel HBA
- Properly installed vendor HBA drivers (Broadcom/Emulex, QLogic, etc.)
- Identify the World Wide Name (WWN) of HBA ports, which need to be zoned in FC switches with storage WWNs.
- Physical FC links to the SAN switches should be active and the WWN of the HBA should be visible on the FC switch.

FC Switch Zoning

Observe the following best practice to create zones on FC switches:

- Single initiator zoning is recommended (one HBA port per zone)
    - One Host HBA port (Initiator) WWN.
    - it can contain one or more VSP storage port/s WWNs.

Create a Zone with a previously identified Host HBA port WWN with a correct storage port WWN.
This zoning model is identical to all standard FC-SCSI environments.

```
Effective configuration:   (Displaying the effective Zone after the zone is created)
 cfg:    Brocade
 zone:  WIN_SCVMM
                10:00:00:10:9b:d1:8e:70
                10:00:00:10:9b:57:9d:8d
                50:06:0e:80:28:27:18:23
```

LUN Provisioning on VSP

On the storage side, configure the following:

- Create an LDEV (logical volume)
- Map the LDEV to a port (port security, host group, LUN ID)
- Add the host's WWNs to the Host Group
- Assign the LDEV to the host via a unique LUN ID

After zoning and mapping are complete, Windows will detect the LUN automatically.

Multipathing (MPIO) in Windows

Installing and Configuring MPIO (GUI method)

This guide details how to install and configure Microsoft Multipath I/O (MPIO) on Windows Server 2025 using the GUI.

© Hitachi Vantara LLC 2026. All Rights Reserved

Microsoft Windows Host Configuration with HV Storage

Setting up a SAN with a Windows host and VSP storage requires several key steps. Start by installing a supported version of the Windows operating system on the host. Verify compatibility using the *Hitachi Product Compatibility Guide* (PCG) at https://compatibility.hitachivantara.com/products/interop-matrix?keywordsearch=win&os=304,474.

Install the Initiator or Host Bus Adapter (HBA) along with the appropriate driver package. For Fibre Channel protocols, confirm the driver package using the PCG for HBA drivers. For iSCSI configurations, set up the respective initiator software and validate connectivity to the target ports.

HBAs are critical in SAN environments, enabling high-speed, reliable data transfer between servers and storage systems. This section focuses on configuring the I/O stack on a Windows system, which can be optimized for performance and reliability. Key tuning parameters include adjusting HBA port timeouts, queue depth, and other performance-related settings. Depending on the HBA model, some adjustments might need to be made through the HBA BIOS. Refer to the vendor's documentation for detailed BIOS-level configuration instructions.

After installation, verify that Fibre Channel adapters and their drivers are functioning correctly. After this is confirmed, identify the WWN of each HBA that will connect to the storage system.

To identify and gather details about FC HBAs on a Windows host using Device Manager:

- Open Device Manager and navigate to **Storage Controllers**.
- Expand the list to view the installed HBAs.
- Right-click the HBA and select **Properties**.
- Verify that supported drivers are installed. See the Hitachi Vantara PCG.

Ensure that the storage system is fully configured before attaching it to the Windows host, following the guidelines provided in the Hitachi storage system configuration documentation.

Target (Hitachi storage) Queue Depth

Queue depth defines the number of I/O operations a storage device (such as a disk or LUN) can process simultaneously. It determines how many read/write requests can be issued before waiting for previous operations to complete. Adjusting queue depth settings correctly helps optimize I/O performance for connected devices.

The ideal queue depth (also referred to as max tag count) should be carefully evaluated based on the operational environment and workload requirements. Additionally, queue depth recommendations vary depending on the storage system model.

Note: Higher queue depth settings (for example, >2,048 per port) on the Initiator/HBA can impact response times. Exercise caution when using queue depth settings that are higher than these values.

See *Host queue depth* in the *Open-Systems Host Attachment Guide for Virtual Storage Platform Family* on the Product Documentation portal (docs.hitachivantara.com) for information about storage queue depth settings.

Initiator Queue Depth Settings

Queue depth values are configured on the Initiator/HBA based on the storage port queue depth values and number of Logical Devices (LDEVs) associated with the host. Queue depth values should be adjusted (on the Initiator/HBA) either on a "per-Initiator" or "per port" basis, or according to the number of LDEVs associated with the host. Consult the HBA (Initiator) manufacturer's reference documentation for specific guidance to achieve best performance.

© Hitachi Vantara LLC 2026. All Rights Reserved

Volume Management on VSP Storage

Configure Port Topology

In the VSP storage system, set the port configuration to **Fabric ON with Point-to-Point (P-P) mode** for SAN environments that use switches, and to **Fabric OFF with Point-to-Point (P-P) mode** for direct connections between the HBA and storage ports. Proper port configuration ensures stable path discovery and optimized throughput between the Windows host and Hitachi storage.

Host Group and LUN path configuration

Windows hosts and the storage system must be physically connected through FC cables with related components, so that I/O paths can be established between the hosts and VSP storage. See Diagram 1.1 and 1.2 for intended configurations. These paths determine which host has access to which logical volume.

Logical volumes accessible to hosts are referred to as logical units (LUs), and the connections between hosts and LUs are referred to as LU paths.

Before LU paths are defined, server hosts must be classified into host groups.

The following diagram illustrates creating a Host Group, assigning LUNs, adding Host WWNs, and setting the correct Host Mode and Host Mode Options as best practices.

Diagram 1.1: Non-Clustered Environment

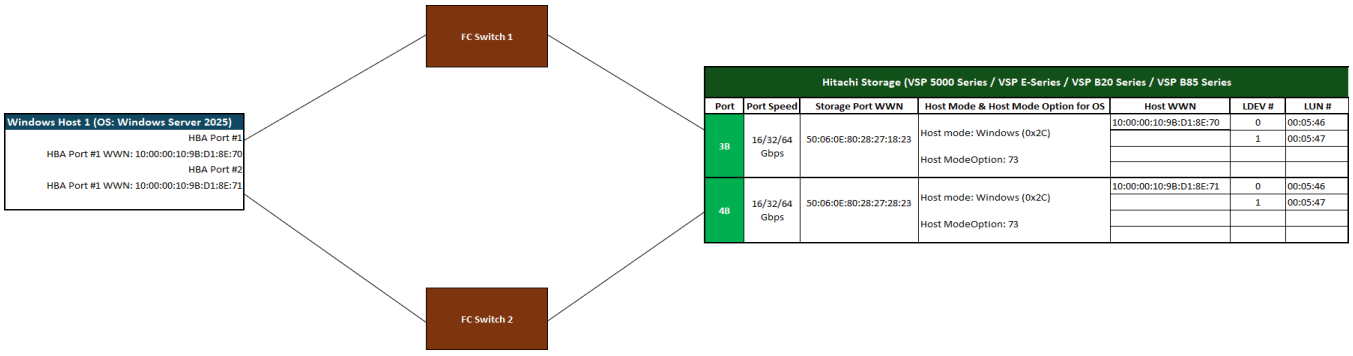

Diagram 1.2: Clustered Environment

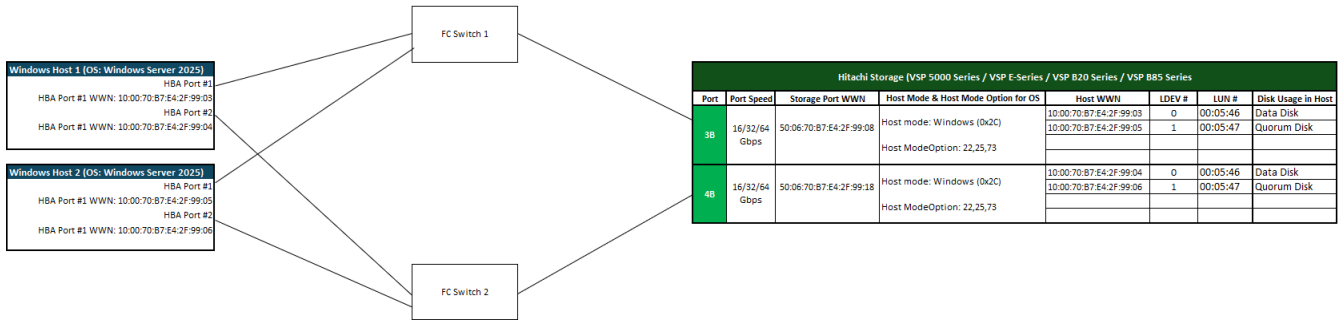

© Hitachi Vantara LLC 2026. All Rights Reserved

Configuring Host mode and Host mode options

On the Hitachi Virtual Storage Platform (VSP) storage system, configure the Host Mode and Host Mode Options (HMOs) specific to Windows to ensure proper communication and command processing between the host and the storage system.

Identify the correct host mode, review the applicable HMOs, and then apply these settings to the storage system.

Carefully review and determine which host modes you will set when configuring your storage system and observe all cautions and restrictions for the host modes. The host mode and host mode options must be set on the host-facing port before the host is connected to that port. If you change the host mode or HMOs after the host is connected, you must reboot the host (server) for the new settings to be recognized. 2C [Windows Extension] is the Host mode for Windows server hosts.

See the Product Documentation portal (docs.hitachivantara.com) for more information:

- *Host modes and host mode options for host-facing host ports* in the *Provisioning Guide for VSP E Series*
- *Host modes and HMOs for SCSI host-facing ports in the Provisioning Guide for VSP One Block 20 Series*
- *Provisioning Guide for Open Systems for VSP 5000 Series*
- *Provisioning Guide for Open Systems for VSP One Block 85*

To protect mission-critical data in your storage system from illegal access, apply security policies to logical volumes. Use LUN Manager to enable LUN security on ports to safeguard LUs from illegal access.

If LUN security is enabled on ports, host groups affect which host can access which LUs. Hosts can access only the LUs associated with the host group to which the hosts belong. Hosts cannot access LUs associated with other host groups. For example, hosts in the Linux host group cannot access LUs associated with the Windows host group. Also, hosts in the Windows host group cannot access LUs associated with the Linux host group.

© Hitachi Vantara LLC 2026. All Rights Reserved

Install MPIO on Windows Server 2025

MPIO is an optional feature in Windows Server 2025 that is not enabled by default. You can install it using the **Add Roles and Features Wizard** in Server Manager.

Install MPIO

- Open Server Manager:
  - Click the **Start** button and select **Server Manager**.
- Launch the Wizard:
  - In the Server Manager dashboard, click **Manage** in the top-right menu bar and select **Add Roles and Features**.
- Navigate to Features:
  - Click **Next** repeatedly to skip through the **Before You Begin**, **Installation Type**, **Server Selection**, and **Server Roles** pages until you reach the **Features** page.
- Select MPIO:
  - Scroll down the list of features, check the box next to **Multipath I/O**, and click **Next**.
- Confirm and Install:
  - On the **Confirmation** page, check the box **Restart the destination server automatically if required** (optional but recommended for MPIO).
  - Click **Install**.
- Complete Installation:
  - After the installation finishes, click **Close**.

Note: You must restart the server to finalize the installation. If you did not select the automatic restart option, manually restart the computer now.

Accessing the MPIO Control Panel

You can access the MPIO configuration settings using either Server Manager or the Windows Start menu.

Method 1: Using Server Manager (Recommended)

1. Launch **Server Manager**.
2. Click the **Tools** menu located in the top-right corner of the dashboard.
3. Select **MPIO** from the drop-down list.

Method 2: Using the Start Menu

1. Click **Start** or press the Windows key.
2. Type MPIOin the search bar.
3. Select the MPIO application from the results to launch the **Properties** window.

Understanding the MPIO Properties Dialog Box

The MPIO Properties window consists of four tabs used to manage multipath configurations.

MPIO Devices Tab

This is the default tab that appears when you open the control panel.

- Function: It lists the Hardware IDs (Vendor and Product strings) of all storage devices currently managed by MPIO.
- How it works: MPIO manages devices by matching their Hardware IDs to a system-maintained list called MPIOSupportedDeviceList. This list is automatically populated by the Device-Specific Module (DSM) information file (.inf) during installation.
- Adding a Device: To manually enable MPIO for a specific storage system or device, click **Add**.

© Hitachi Vantara LLC 2026. All Rights Reserved

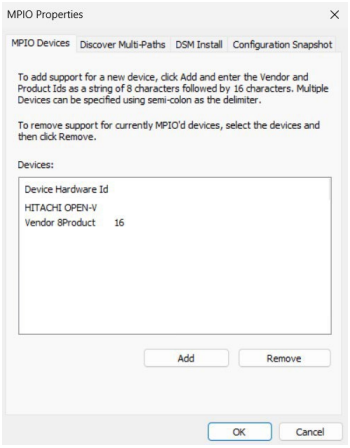

Diagram 2

Configuring the parameter (Microsoft DSM/MPIO)

Set the value of **PDORemovePeriod** in the registry 20.
See *Restrictions on the configuration for connecting hosts and storage systems* in Chapter 4: Managing logical volumes of the *Provisioning Guide for Open Systems for VSP 5000 Series* on the Product Documentation portal (docs.hitachivantara.com).

Procedure

1. Start Windows PowerShell.
2. Use the **Get-MPIOSetting** command to check the value of PDORemovePeriod.

```
PS C:\Users\Administorator> Get-MPIOSetting
PathVerificationState : Enabled
PathVerificationPeriod : 30
PDORemovePeriod : 20
RetryCount : 3
RetryInterval : 1
UseCustomPathRecoveryTime : Enabled
CustomPathRecoveryTime : 15
DiskTimeoutValue : 60
```

3. Run the following command if required.

```
Set-MPIOSetting -NewpathVerificationState enabled
Set-MPIOSetting -CustomPathRecovery enabled
Set-MPIOSetting -NewPathRecoveryInterval 15
Set-MPIOSetting -NewPDORemovePeriod 20
```

   Use the **Get-MPIOSetting** command to verify that the value of all required parameters.

4. Restart the host OS.

© Hitachi Vantara LLC 2026. All Rights Reserved

Windows Failover Cluster

Windows Failover Clustering (WFC) provides high availability for applications and services by grouping multiple servers (nodes) that work together. If one node fails, the workload automatically moves to another node.

Quorum Configuration

The Quorum configuration determines the number of node failures the cluster can sustain while remaining online.

Quorum Modes: A quorum mode is configured at the WSFC cluster level that dictates the methodology used for quorum voting. The Failover Cluster Manager utility will recommend a quorum mode based on the number of nodes in the cluster. The following quorum modes can be used to determine what constitutes a quorum of votes:

- Node Majority
- Node and File Share Majority
- Node and Disk Majority
- Disk Only

Best Practice: Use “Disk Only” to ensure voting capability. A shared disk in cluster resource is designated as a witness which can be accessed by all the cluster nodes. Use a shared LDEV from Hitachi storage for the best experience and reliability.

See diagram 1.2 for a Quorum Disk example.

Node Failover and Failback Policy

Failover settings control how resources (roles/VMs) move during a failure, while failback settings control how they return when the primary node recovers.

Failover Settings

- Preferred Owners: Explicitly set the order in which nodes should host a specific role. For a 4-node cluster hosting 4 SQL instances, assign a different primary owner for each instance to balance the load.
- Possible Owners: Ensure all nodes that can technically run the workload are checked. Uncheck nodes only if they lack the hardware/connectivity needed for that specific role.

Failback Settings

Failback occurs when a preferred node comes back online after a failure.

- Best Practice: Set Failback to "Allow Failback between [Hour] and [Hour]" (for example., 20:00 to 06:00).
- Why: Immediate failback can cause "flapping" (a service moving back and forth rapidly if a node is unstable), causing double downtime. Configuring a maintenance window ensures the node is stable before the workload returns.

Network Configuration

- Heartbeat (Cluster Traffic)
- Public (Management Traffic, optionally can be Management Traffic and Cluster Traffic for added redundancy for Cluster)

Correct network isolation ensures that heavy client traffic does not block critical cluster communication (heartbeats).

© Hitachi Vantara LLC 2026. All Rights Reserved

Diagram 3

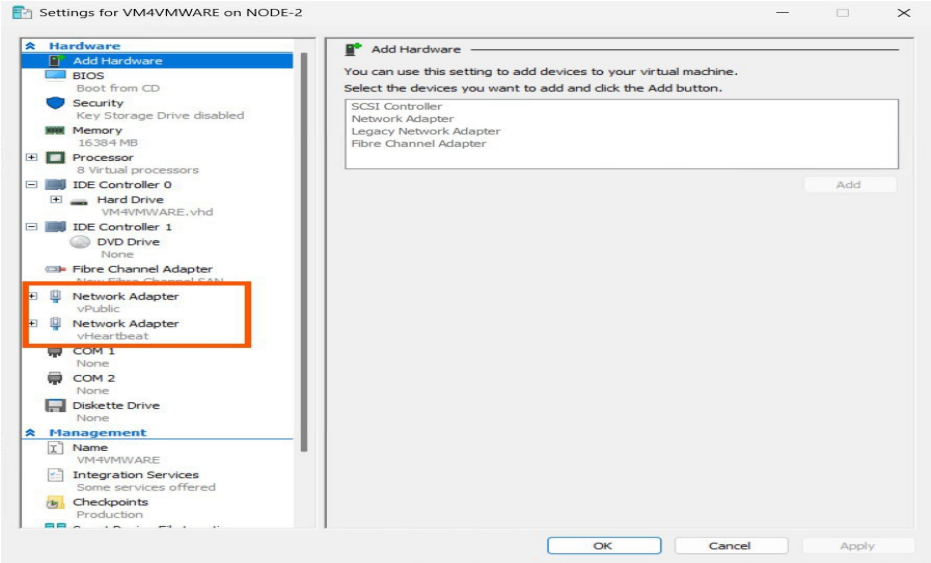

Virtual Heartbeat Network (Cluster Only)
Purpose: Used for inter-node health checks and CSV redirection traffic.
Cluster Use: Configure as "Cluster Only" or "Internal".

- Configuration:
  - Subnet: It should be a different subnet from the public network.
  - DNS: Register this connection in DNS.
  - Assign static IP addresses for each network interface (Private).
  - Ensure correct subnetting and DNS settings.

Virtual Public Network (Client and Cluster)
Purpose: Used for client applications to connect to the cluster and for management.
Cluster Use: Configure as "Cluster and Client".

- Configuration:
  - DNS: Register this connection in DNS.
  - Default Gateway: Configured.
  - Assign static IP addresses for each network interface (Public).
  - Ensure correct subnetting and DNS settings.

Multipath I/O (MPIO): Configure all production hosts with at least two data paths to the VSP storage system to eliminate single points of failure. Single-path I/O should be restricted to non-critical test or development environments.

Validating the Cluster Configuration

Before deploying a production Hyper-V cluster on Hitachi storage, run a full cluster validation. All storage and MPIO tests must pass to ensure stability.

1. Configure Connectivity: Establish two or more data paths from each host to the VSP One Block storage.
2. Install MPIO: Install and configure the MPIO feature on every BareMetal node.
3. Create Host Groups: Create a Host Group on VSP Storage for the cluster nodes.
4. Map Volumes: Map shared cluster volumes (LUN) to the Host Group using a consistent LUN ID across all hosts.
5. Initialize Disk: On one host, initialize the new disk, bring it online, and format it as NTFS.
6. Rescan Disks: Perform a disk rescan on all other hosts in the cluster to ensure visibility.
7. Run Validation: Use Failover Cluster Manager to run the "Validate Configuration" wizard for the hosts.
8. Verify Results: Confirm that all tests, specifically those related to Disk and MPIO, pass successfully.
9. Troubleshoot: If any tests fail, the configuration is not supported for clustering. Resolve the specific disk or MPIO errors and re-run the validation until all tests pass.

© Hitachi Vantara LLC 2026. All Rights Reserved

Cluster Shared Volumes (CSV)

Format the disks designated for Cluster Shared Disk with NTFS Filesystem.

The Add-ClusterSharedVolumecmdlet makes a volume available in the Cluster Shared Volumes in a failover cluster. The Cluster Shared Volume must be chosen from the disks in Available Storage.

This example adds Cluster Disk 4 to the Cluster Shared Volumes on the cluster:

```
Add-ClusterSharedVolume -Name "Cluster Disk 4"
```

© Hitachi Vantara LLC 2026. All Rights Reserved

Windows Hyper-V on Failover Cluster

Hyper-V Best Practices on VSP Block Storage

Microsoft Hyper-V leverages Windows Failover Clustering to deliver a resilient, high-availability virtualization platform where workloads remain operational despite hardware failures. This architecture relies heavily on shared VSP One Block storage – accessed via Fibre Channel. To enable Live Migration, Cluster Shared Volumes (CSV) are necessary. By offering low-latency performance, seamless scalability, and robust fault tolerance, VSP One Block storage ensures that multiple cluster nodes can access data simultaneously, making it the ideal foundation for supporting mission-critical virtual machines and high-demand applications.

Recommended Configuration Guidelines

The following best practices apply to any Hyper-V deployment on VSP storage:

- Windows Server Datacenter (Desktop Experience): Utilize the Datacenter edition with Desktop Experience to provide a full graphical user interface (GUI), which simplifies management and ensures compatibility with tools requiring a local desktop environment.
- Updates and Security: Ensure all hosts and VMs are patched regularly and have adequate malware protection.
- Resource Optimization: Disable unused hardware devices and services to free up host resources and reduce power consumption.

Hyper-V Networking Fundamentals

- Hyper-V networking is built on two core components: the virtual switch and the virtual network adapter. At least one virtual switch and one virtual network adapter are needed to provide network connectivity to a virtual machine (VM).
- The virtual switch enables connectivity to Ethernet-based networks, while the virtual network adapter connects a VM to a port on the virtual switch, allowing the VM to communicate with other systems on the network.

Basic Network Configuration

- The simplest approach to configuring networking in Hyper-V is to create a virtual switch during Hyper-V installation. When a VM is created and connected to this switch, Hyper-V automatically assigns a virtual network adapter to the VM.
- Additional virtual switches and network adapters can be created to support different networking requirements. Virtual switches are host-level components and can be shared by multiple VMs, whereas each virtual network adapter is associated with a single VM.

Types of Hyper-V Virtual Switches

Hyper-V provides three types of virtual switches to support different networking scenarios:

External Virtual Switch

- Binds to a physical network adapter
- Provides connectivity to the physical network
- Allows VMs and the host to communicate with external systems
- Recommended for production workloads

Internal Virtual Switch

- Enables communication between VMs on the same host
- Allows communication between the host and VMs
- Does not provide access to external networks

Private Virtual Switch

- Enables communication only between VMs on the same host
- No connectivity to the host or external networks

© Hitachi Vantara LLC 2026. All Rights Reserved

- Typically used for isolated or test environments

Best Practice for Clustered and Non-Clustered Environments

1. **Virtual Switch Design**
   - Use External vSwitch for VM workloads
   - Bind to physical NICs (pNICs).
   - Do not use private/internal switches for production VMs.
   - One vSwitch per physical NIC team
2. **Non-Clustered Hosts**

Use separate NICs or VLANs for:

- VM traffic
- Management (Host OS)
- Storage/iSCSI/SMB
- Backup traffic
- Live Migration (if enabled)

Note the following:

- VLAN separation is fine for lab/small deployments.
- Physical separation preferred for production.

3. **Network Layout for Clustered Hyper-V (Failover Clustering)**

A typical production cluster should segment:

- Management Network
  - Host OS, AD, domain traffic
  - Enable only required protocols
  - VM traffic (not allowed)
- Cluster Heartbeat Network
  - Ensures nodes stay in sync
  - Low latency
  - No routing, no blocking
  - Often combined with CSV traffic in modern deployments
- Live Migration Network
  - Dedicated NIC or VLAN
  - Use SMB or compression mode
  - Prefer high-speed NICs
- VM Network(s)
  - One or more external vSwitches
  - Based on workload segregation
  - Often delivered via converged networking

4. **Converged Networking Recommendations**

Converged networking = a few high speed NICs carrying multiple traffic types via vNICs. Best practice for production (clustered)

Use 2 × 25/40/100GbE NICs with:

- SET team
- vSwitch on top of SET
- vNICs created for:
  - Management
  - Cluster
  - CSV
  - Live Migration
  - Backup
  - Others as needed
- Bandwidth Management
  - Use QoS to guarantee:
  - CSV/Storage traffic minimum bandwidth
  - Live Migration max cap to avoid congestion.

Processor and Memory Allocation

Allocate vCPU and vMemory resources according to the application's performance needs and the importance of its role within the enterprise.

For Windows Server 2025 virtual machines (VMs), the vCPU and vMemory allocations are: vCPU

© Hitachi Vantara LLC 2026. All Rights Reserved

- Minimum: 1 vCPU (A 1.4 GHz 64-bit processor is the minimum physical requirement, and a single core is the minimum allocation, though performance will be limited).
- Maximum: 2,048 vCPUs (for a Generation 2 VM using Hyper-V).

vMemory

- Minimum: 800 MB (for a Server Core installation in a VM; attempting to install with the bare minimum of 512 MB might fail).
- Maximum: 240 TB (terabytes) for a Generation 2 VM.

Best Practices for Choosing Hyper-V Virtual Machine Generation

The choice between creating a Generation 1 or Generation 2 virtual machine depends on the guest operating system to be deployed and the boot method. In general, Generation 2 virtual machines are recommended to take advantage of modern features such as UEFI firmware and Secure Boot.

However, a Generation 1 virtual machine should be selected in the following scenarios:

- An existing prebuilt virtual hard disk (VHD or VHDX) is being used that is not compatible with UEFI
- The target operating system is not supported by Generation 2 virtual machines
- The required boot method is not supported in Generation 2

For a detailed comparison of available features and supported guest operating systems, see the Hyper-V feature compatibility by generation and guest documentation.

Note that the virtual machine generation cannot be changed after creation. Therefore, it is strongly recommended to carefully evaluate the operating system requirements, boot options, and feature dependencies before selecting the virtual machine generation.

Best Practice: Choose a Generation 2 Virtual Machine for operating systems such as Windows Server 2025 and Windows Server 2022.

Virtual Machines on Hitachi SAN Disks using Virtual SAN

Deploying Hyper-V virtual machines on Hitachi SAN storage provides enterprise-grade performance, availability, and scalability. To achieve optimal results, storage, host, and Hyper-V configurations must align with recommended best practices.

Virtual Fibre Channel (Virtual SAN)

Hyper-V Virtual Fibre Channel (vFC) is a mechanism that allows a Virtual Machine (VM) to bypass the host OS storage stack and communicate directly with the VSP storage system. This presents the LUNs directly to the Guest OS, enabling "Guest Clustering" (for example, SQL Server Always On Failover Cluster Instances) and high-performance workloads that require raw disk access.

vFC relies on **N_Port ID Virtualization (NPIV)**. NPIV allows a single physical Fibre Channel HBA port to register multiple unique World Wide Port Names (WWPNs) with the SAN fabric. To the SAN switch and the VSP storage, the VM appears as a distinct physical endpoint, independent of the Hyper-V host.

Prerequisites:

- Host Hardware: Physical HBAs must support NPIV and have drivers that expose this capability to Hyper-V.
- Fabric Config: NPIV must be enabled on the SAN switch ports connected to the Hyper-V hosts.
- Topology: To maintain redundancy, you must configure two Virtual SANs in Hyper-V:
    - Virtual SAN A Mapped to Physical second HBA Port 1 (Fabric A)
    - Virtual SAN B Mapped to Physical second HBA Port 2 (Fabric B)

Configuration Steps

Setting up vFC involves configuring the host, the VM, and the guest OS.

© Hitachi Vantara LLC 2026. All Rights Reserved

Step 1: Define the Virtual SAN on the host
Use the Virtual SAN Manager in Hyper-V to bind a physical HBA port to a logical "Virtual SAN."

- *Action:* Create one Virtual SAN for each physical fabric (for example, "vSAN-Fabric-A" and "vSAN-Fabric-B").

Configure Hyper-V SAN Manager
New Virtual SAN Created with the Name ”New Fibre Channel SAN”
Diagram 4.1

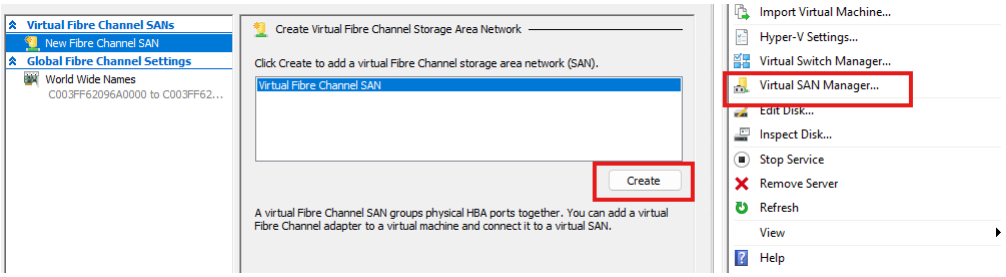

Diagram 4.2

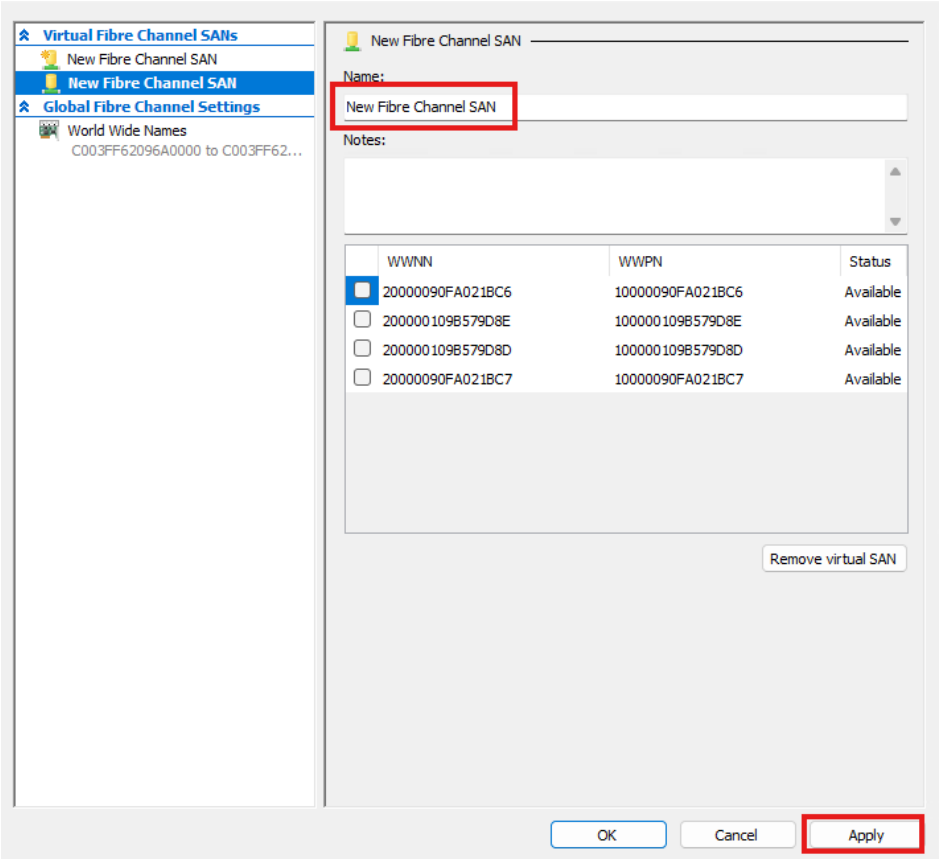

Step 2: Add vFC-HBA to the VM - from inside the VM settings, add a "Fibre Channel Adapter."

- *Best Practice:* Add two adapters per VM. Connect Adapter 1 to "vSAN-Fabric-A" and Adapter 2 to "vSAN-Fabric-B" to ensure multipathing.
- Add Fibre Channel Adapters (NPIV) to connect to the shared storage LUNs.

Note: To add 2 x vFC-HBAs, 2 x Virtual Fibre Channel SANs must be created.

WWN Configuration (Set A / Set B)

Live Migration moves a running VM from one physical host to another without downtime. To support this with Fibre Channel, Hyper-V assigns two sets of WWNs to each virtual HBA:

- Set A: The active addresses used on the current host.
- Set B: The standby addresses used for the destination host.

© Hitachi Vantara LLC 2026. All Rights Reserved

When creating the Host Group on VSP storage, you must add all four WWNs for a single VM (2 vHBAs x 2 sets) to the Host Group.

Diagram 5

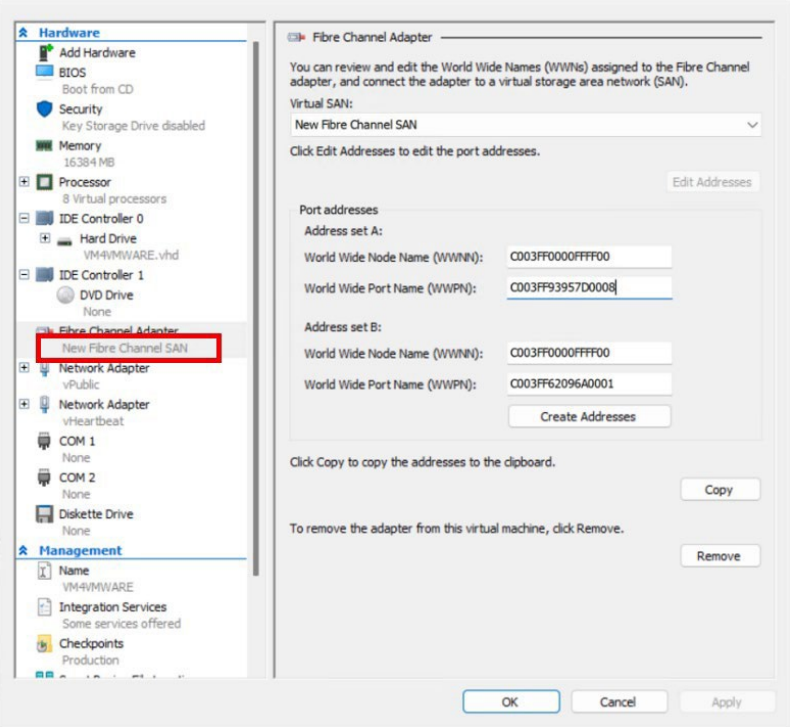

Fibre channel switch section

- The switch port will show NPIV Login.

```
SW6510:admin> nsshow
{
 Type Pid    COS     PortName                NodeName                 TTL(sec)
 N    500800;3;10:00:00:10:9b:d1:8e:6f;20:00:00:10:9b:d1:8e:6f; na
    FC4s: FCP
    PortSymb: [34] "Emulex PPN-10:00:00:10:9B:D1:8E:6F"
    NodeSymb: [74] "Emulex LPe35002-M2 FV14.2.539.16 DV12.6.165.8000 HN:NODE-1 OS:Windows
2019"
    Fabric Port Name: 20:08:00:27:f8:43:7d:78
    Permanent Port Name: 10:00:00:10:9b:d1:8e:6f
    Port Index: 8
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 16G
 N    500900;

>>>>>Truncated<<<<<<

3;c0:03:ff:93:95:7d:00:08;c0:03:ff:00:00:ff:ff:00; na
    FC4s: FCP
    PortSymb: [64] "Emulex PPN-C0:03:FF:93:95:7D:00:08 Vport-1 VName-Hyper-V VM Port"
    NodeSymb: [74] "Emulex LPe35002-M2 FV14.0.505.39 DV12.6.165.8000 HN:NODE-2 OS:Windows
2019"
    Fabric Port Name: 20:15:00:27:f8:43:7d:78
    Permanent Port Name: 10:00:00:10:9b:57:9d:8e
    Port Index: 21
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
```

© Hitachi Vantara LLC 2026. All Rights Reserved

```
    Device link speed: 16G
 N    501503;    2,3;c0:03:ff:93:95:7d:00:0a;c0:03:ff:00:00:ff:ff:00; na
    FC4s: FCP
    PortSymb: [64] "Emulex PPN-C0:03:FF:93:95:7D:00:0A Vport-2 VName-Hyper-V VM Port"
    NodeSymb: [74] "Emulex LPe35002-M2 FV14.0.505.39 DV12.6.165.8000 HN:NODE-2 OS:Windows
2019"
    Fabric Port Name: 20:15:00:27:f8:43:7d:78
    Permanent Port Name: 10:00:00:10:9b:57:9d:8e
    Port Index: 21
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 16G
The Local Name Server has 11 entries }

>>>>>Truncated<<<<<<
```

- Switch Zoning:
  Make zone in the switch and add all host (vWWN) and storage WWNs for a Virtualized Windows Server Failover Cluster (vWSFC).

```
zone:  WIN_SCVMM1
                50:06:0e:80:28:27:18:23  Storage Port WWN
                c0:03:ff:93:95:7d:00:08  vFC (Set-A) Port WWN
zone: WIN_SCVMM2
                50:06:0e:80:28:27:18:23  Storage Port WWN (Same Storage port WWN)
                c0:03:ff:93:95:7d:00:0a   vFC (Set-B) Port WWN
```

Recommended Approach

- Use Cluster Shared Volumes (CSV) for VM storage in clustered environments

LUN Provisioning to Cluster

Create Host Group for VM using their WWPNs and subsequently map the LUNs to the Host Group with the correct WWN of the virtual machines.

Assigning the vSAN WWPN:

```
raidcom add hba_wwn -port CL3-D-1 -hba_wwn c003ff93957d0008 -IH6811 raidcom add hba_wwn -port CL4-D-1 -hba_wwn
c003ff93957d000a -IH6811
raidcom get hba_wwn -port CL3-D-1 -IH6811
PORT GID GROUP_NAME           HWWN Serial# NICK_NAME CL3-D              1 WIN    c003ff93957d0008 810008 -

raidcom get hba_wwn -port CL4-D-1 -IH6811
PORT GID GROUP_NAME   HWWN Serial# NICK_NAME CL4-D              1 WIN    c003ff93957d000a 810008 -
```

Adding Luns to the Host Group:

```
raidcom add lun -port CL3-D-1 -lun_id 0 -ldev_id 05:46 -IH6811 raidcom add lun -port CL4-D-1 -lun_id 1 -ldev_id
05:47 -IH6811
```

© Hitachi Vantara LLC 2026. All Rights Reserved

Terms and definitions

The following terms are commonly used with virtual Fibre Channel:

- **MPIO** Multipath I/O (MPIO) software is used to provide load balancing and fault tolerance for multiple virtual Fibre Channel connections in the same server running Hyper-V.
- **NPIV** Virtual Fibre Channel for Hyper-V guest operating systems uses the existing N_Port ID Virtualization (NPIV) T11 standard to map multiple virtual N_Port IDs to a single physical Fibre Channel N_port. A new NPIV port is created on the server running Hyper-V each time you start a virtual machine that is configured with a virtual HBA. When the virtual machine stops on the Hyper-V server, the NPIV port is removed.
- **Virtual SAN** A virtual SAN (vSAN) defines a named group of physical Fibre Channel ports that are connected to the same physical SAN.
- **World Wide Names** A World Wide Port Name (WWPN) is a unique number that is provided to a virtual Fibre Channel HBA, similar to a MAC address. This unique key is used to allow the storage fabric to recognize a specific HBA. A World Wide Node Name (WWNN) is assigned to a node (an endpoint, a device) in a virtual Fibre Channel fabric.

© Hitachi Vantara LLC 2026. All Rights Reserved